

SDSRV Design Status

Jan Dreisbach

jdreisba@eos.hitc.com

18 April 1996

706-CD-003-001 Day 4 Book A

JD2-1

Agenda



- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- ESDT Design
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics
- Next Steps
- Design Topics

706-CD-003-001 Day 4 Book A

JD2-2

Release B Capabilities



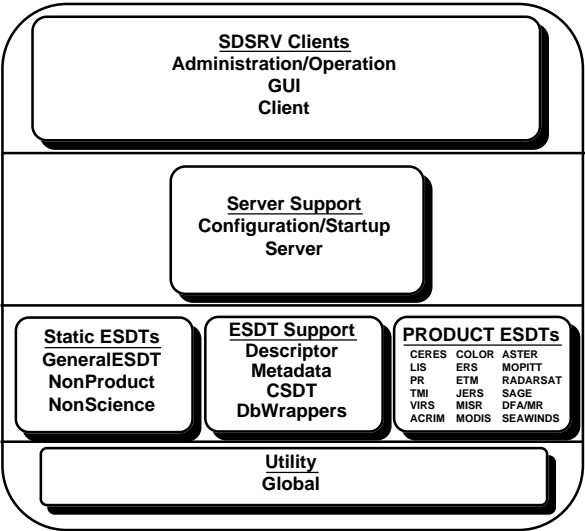
ESDT-Related Work

- Multiple Command Request
- New ESDTs - 115 (Total of 173)
- On-demand Processing
- Subsetting

SDSRV Internal Infrastructure

- Suspend and Resume Sessions
- Accounting
- Metadata Problem Reports
- ASTER DAR Support
- ESQL Support
- Migration from Sybase to Illustra

Software CSCs



Software CSC Status



SDSRV Internal Infrastructure CSCs

CSC	Abbrev.	"A" Unaffected	Augmented From "A"	Exclusively for "B"
Administration/Op	Ad		Low	
Client	Cl		8 New Attributes, 31 New Operations	
Configuration/Startup	Cn	X		
Metadata	Md		9 New Operations	
CSDT	Cs		1 New Class, 21 New Operations	
DB Wrappers	Db		High	
Descriptors	De		1 New Operation	
General ESDT	Gb		1 New Attribute, 5 New Operations	
Global	Gl		Low	
GUI	Gu		Low	
Server	Sr		10 New Classes, 6 New Attributes, 38 New Operations	

706-CD-003-001 Day 4 Book A

JD2-5

Software CSC Status (cont.)



ESDT Data Type CSCs

CSC	Abbrev.	"A" Unaffected	Augmented From "A"	Exclusively for "B"
CERES	Ce		Low	
LIS	Li		Low	
ESDTs	Np		Low	
Non-Science ESDTs	Ns		2 New Classes	
PR	Pr		Low	
TMI	Tm		Low	
VIRS	Vi		Low	
ACRIM	Ac			1 New Class
ASTER	As			1 New Class
COLOR	Cc			1 New Class
ERS	Er			1 New Class
ETM	Et			1 New Class
JERS	Je			1 New Class
MISR	Mi			1 New Class
MODIS	Mo			1 New Class
MOPITT	Mp			1 New Class
RADARSAT	Ra			1 New Class
SAGE	Sa			1 New Class
SSA	Ss			1 New Class
SeaWinds	Sw			1 New Class

706-CD-003-001 Day 4 Book A

JD2-6

Agenda



- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- ESDT Design
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics
- Next Steps
- Design Topics

706-CD-003-001 Day 4 Book A

JD2-7

Progress Since IDR-B



Overall SDSRV Progress:

- Closed all IDR-B RIDS
- Held RID #5 Meeting - 2/96
- Supported Infrastructure Implementation Feasibility Workshop - 2/96
- Supported OPS Workshop - 2/96
- Held ESDIS/DAAC Issue Resolution Mtg - 2/96
- Held ESDIS TIM - 2/96
- Worked with Rel A to minimize design breakage for Rel B
- Refined technical concept of operations for key design areas
- Provided input to User Modeling effort

Completed Input to CDR-B DID 313-CD-006-002 Delivery:

- Finalized public interfaces with other subsystems

706-CD-003-001 Day 4 Book A

JD2-8

Progress Since IDR-B (cont.)



Completed Input to CDR-B DID 604-CD-002-003 and DID 605-CD-002-001 Deliveries:

- Finalized operational scenarios

Completed and Delivered CDR-B DAAC Specific Design Documents:

- DID 305-CD-31-002
- DID 305-CD-33-002
- DID 305-CD-34-002
- DID 305-CD-35-002
- DID 305-CD-36-002
- DID 305-CD-37-002
- DID 305-CD-38-002
- DID 305-CD-39-002

706-CD-003-001 Day 4 Book A

JD2-9

Progress Since IDR-B (cont.)



ESDT-Related Progress:

- Maintained 304-CD-005-002 Appendix F (Data Types Matrix) based on feedback from Instrument Teams
- Evaluated HDF-EOS Specification based on “B” requirements
- Validated Rel A ESDT architecture and life-cycle approach

On-Demand Services Progress:

- Resolved On-Demand processing with PDPS by completing interfaces and protocol definition

Subsetting Progress:

- Incorporated SSIG proceedings received 3/96 into technical concept of operations

706-CD-003-001 Day 4 Book A

JD2-10

Progress Since IDR-B (cont.)



Accounting Progress:

- Completed design via “shared” use of Data Distribution Utilization CSC (Uz) classes and methods

ASTER DAR Progress:

- Completed definition of SDSRV role in ASTER DAR support (with Data Engineering) that uses existing class constructs

Sybase to Illustra Migration Progress:

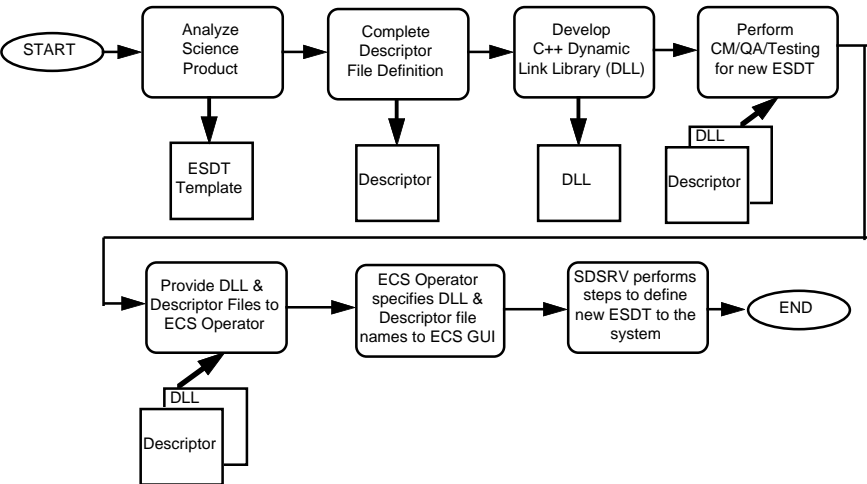
- Completed significant portion of Sybase/Illustra Benchmark
- Completed Illustra SOW and began work on Illustra Manageability SOW
- Completed Illustra trades - CDR-B

Agenda



- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- **ESDT Design**
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics
- Next Steps
- Design Topics

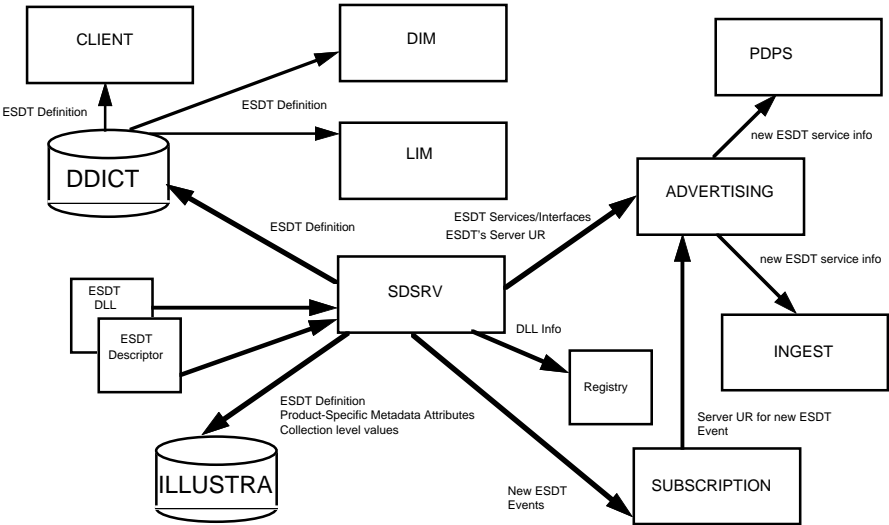
ESDT Lifecycle



706-CD-003-001 Day 4 Book A

JD2-13

ESDT Architecture



706-CD-003-001 Day 4 Book A

JD2-14

ESDT A->B Migration



Rel A ESDTs/Services: 38, 108

Rel B ESDTs/Services: 173, 771

ESDT “Architecture Impact” for Release B

- Release A provides architectural foundation for “generic” ESDT definition
- Release B brings new subsystem interface to Data Dictionary (DDICT)
- Schedule driven decision not to integrate Illustra features into ESDT definition

ESDT A->B Migration (cont.)



ESDT Type Definition for Release B

- Effort required for “analysis” of new data types (Descriptor File Definition)
 - Data Types defined in Release A and B Matrix (See Handout #1)
 - Template developed for Release A provides process for performing analysis (See Handout #2)
- Effort required for code development of services for new data types (DLL)

Evolutionary Enhancements (new features vs. new technology)

- Operational/Administration Tools

Agenda



- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- ESDT Design
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics
- Next Steps
- Design Topics

706-CD-003-001 Day 4 Book A

JD2-17

Schema Trade



Approaches:

- (1) Use “pure” relational schema (Rel A Sybase + Rel B enhancements)
 - Requires minimal labor for (development) transition (+)
 - Minimizes impact to V0 transition (+)
 - A->B transition labor is eased (+)
 - Easier to revert back to Sybase, if necessary (+)
 - Probably lose on performance (-)
 - No positioning for evolution (-)
- (2) Use combination of Illustra ORDBMS features w/ relational schema
 - Reuse schema changes based on benchmark prototype (+)
 - Take advantage of potential performance gains (+)
 - Begin positioning for evolution (+)
 - Transition tasks require mapping Sybase & Illustra schemas (-)
 - During the conversion of stored procedures, more development work required (-)

+ = Advantage
- = Disadvantage

706-CD-003-001 Day 4 Book A

JD2-18

Schema Trade (cont.)



- (3) Maximize use of Illustra ORDBMS features, including approach to ESDTs
 - Positions the SDSRV for maximal use of COTS(+)
 - Causes rework of SDSRV code outside the DBMS Wrappers (-)
 - All benefits/disadvantages of previous option

Selected approach is to use a combination of Illustra ORDBMS features with a relational schema (Option 2).

DBMS Wrapper Trade



- (1) Use current approach with “traditional” C-style DBMS application program interface (API) through Illustra’s LIBMI
 - retains application insulation (+)
 - minimizes labor required to make the Sybase to Illustra transition (+)
 - does not take advantage of COTS evolvability (-)
- (2) Change DBMS Wrapper approach to conform to COM model style of API using the Illustra C++ API
 - significant rework of Rel A code (-)
 - impacts SDSRV code outside the wrappers (-)
 - less code to maintain in the long run (+)
 - easier error handling (+)
 - more “object-oriented” than traditional C-style interface (+)

Selected approach is to use “traditional” C-style DBMS API through Illustra’s LIBMI.

Agenda



- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- ESDT Design
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics
- Next Steps
- Design Topics

706-CD-003-001 Day 4 Book A

JD2-21

Sybase/Illustra Benchmark



Goals

Overview

Findings

706-CD-003-001 Day 4 Book A

JD2-22

Sybase/Illustra Benchmark Goals



Primary

- To mitigate development/production risks by gaining early software, administration, and capacity experience with both Sybase SQS and Illustra products.
- To establish baseline performance measurements for Sybase SQS and Illustra running 2-D spatial queries against ~100 GB of metadata under both light and heavy (~100 users) load conditions.

Secondary

- To establish baseline measurements for metadata (initial) load times in both Sybase and Illustra.
- To establish baseline performance measurements for insert/update/delete of (single) rows of metadata in both Sybase and Illustra under both light and heavy load conditions.

Benchmark Configuration



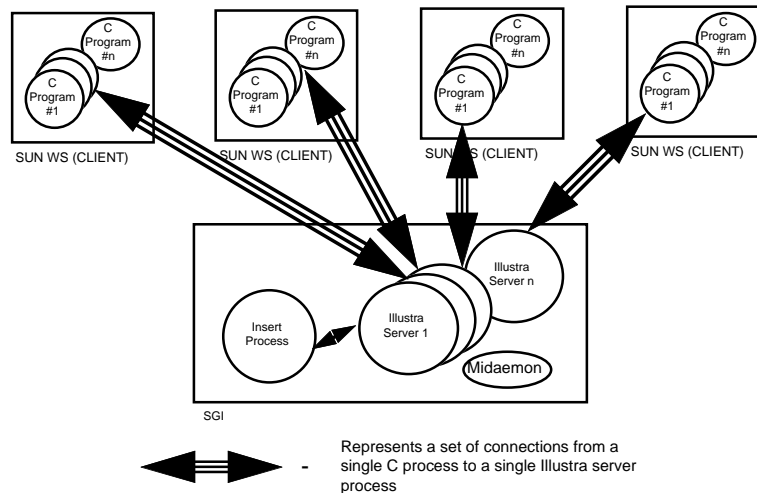
Hardware - SGI Challenge L

- 230 GB of RAID disk
- 512 MB
- Located in the EDF
- Connectivity to vendor machines

Software - IRIX 5.3

- Sybase SQS 2.1/Sybase 10.0.2
 - C/Open Client
- Illustra 3.2.0.21
 - C/LIBMI

Benchmark Architecture



706-CD-003-001 Day 4 Book A

JD2-25

Benchmark Transactions



2-D Spatial Queries

- Scenarios generated by Science Office
- Example: "User queries ECS about the availability of the following data for the Little Washita watershed, OK-soil maps, DEM, Insitu rainfall data, Run-off data.
 - Little Washita watershed coordinates: 36N 98W, 36N 100W, 34N 100W, 34N 98W, 35N 99W
 - Selected product instances will be ASTER07B

2-D Spatial Queries with "attributes"

Metadata Insertion Based on "Sample Ingest"

706-CD-003-001 Day 4 Book A

JD2-26

Benchmark Metadata



Goal:

Generate ~100GB of metadata

- Roughly equal distribution of metadata around the world, i.e., no focus in a particular area of the world
- Include all instruments that generate products that describe a 2-D spatial area
 - MODIS
 - SAGE
 - ASTER
 - RADAR

Benchmark Metadata (cont.)



Approach:

- Use the “Karl Cox Orbit Model” to generate the instrument data
 - Period of observation from 8/17/1998 through 10/1/2002
 - Generating about 300 bytes/granule for prototype
 - “Padding” the remaining 1700 bytes/granule
 - Prototyping with 46,424,701 granules
- Store granule information only - NOT collection data

Benchmark Metadata Composition



ASTER Products

PRODUCT Name	NO. of DB Rows
AST_L1A	1184355
AST_L1B	473742
AST_04	94748
AST_05	94748
AST_06A	473742
AST_06B	473742
AST_06C	473742
AST_07A	94748
AST_07B	94748
AST_08	94748
AST_09A	94748
AST_09B	94748
AST_09C	94748
AST_10	473742
Sub Total for ASTER	4311049

Benchmark Metadata Composition



MODIS AM Products

PRODUCT Name	NO. of DB Rows	PRODUCT Name	NO. of DB Rows
MOD_AM01_L1A	867436	MOD_AM27_L3_DY	536134
MOD_AM02_L1B	867436	MOD_AM27_L3_WK	76539
MOD_AM03_L1A	867436	MOD_AM28_L3_DY	536134
MOD_AM04_L2	394363	MOD_AM28_L3_WK	76539
MOD_AM04_L3_DY	536135	MOD_AM29_L3_DY	76539
MOD_AM04_L3_WK	76539	MOD_AM30_L2	867436
MOD_AM05_L2	394363	MOD_AM33_L3_7DY	76539
MOD_AM06_L2	867436	MOD_AM34_L3_10DY	52330
MOD_AM07_L2	867436	MOD_AM35_L2	867436
MOD_AM08_L2	867436	MOD_AM38_L2	867436
MOD_AM09_L3_BRDF_L3_16DY	33463	MOD_AM40_L3_DY	536134
MOD_AM10_L3_DY	536134	MOD_AMOCCLR_L3_DY	536134
MOD_AM11_L3_30DY	17443	MOD_AMOCCLR_L3_WK	76539
MOD_AM11_L3_7DY	76539	SUB10_MOD_AM09_L3_L2	394363
MOD_AM14_L3_10DY	52330	SUB14_MOD_AM14_L2	867436
MOD_AM14_L3_DY	536134	SUB18_MOD_AM41_L2	394363
MOD_AM15_L4_10DY	52330	SUB25_MOD_AM10_L2	394363
MOD_AM17_L4_10DY	52330	SUB28_MOD_AM29_L2	394363
MOD_AM27_L3_DY	536134	Sub Total for MODIS AM	15589476

Benchmark Metadata Composition



MODIS PM Products

PRODUCT Name	NO. of DB Rows
MOD_PM01_L1A	350759
MOD_PM02_L1B	350759
MOD_PM03_L1A	350759
MOD_PM04_L2	159937
MOD_PM04_L3_DY	216803
MOD_PM04_L3_WK	30971
MOD_PM05_L2	159937
MOD_PM06_L2	350759
MOD_PM07_L2	350759
MOD_PM08_L2	350759
MOD_PM09_L3_BRDF_L3_16DY	13527
MOD_PM10_L3_DY	216803
MOD_PM11_L3_30DY	7119
MOD_PM11_L3_7DY	30971
MOD_PM14_L3_10DY	21359
MOD_PM14_L3_DY	216803
MOD_PM15_L4_10DY	21359
MOD_PM17_L4_10DY	21359

PRODUCT Name	NO. of DB Rows
MOD_PM27_L3_DY	216803
MOD_PM27_L3_WK	30971
MOD_PM28_L3_DY	216803
MOD_PM28_L3_WK	30971
MOD_PM29_L3_DY	216803
MOD_PM30_L2	350759
MOD_PM33_L3_7DY	30971
MOD_PM34_L3_10DY	21359
MOD_PM35_L2	350759
MOD_PM38_L2	350759
MOD_PM40_L3_DY	216803
MOD_PMOCCLR_L3_DY	216803
MOD_PMOCCLR_L3_WK	30971
SUB10_MOD_PM09_13_L2	159937
SUB14_MOD_PM14_L2	350759
SUB18_MOD_PM41_L2	159937
SUB25_MOD_PM10_L2	159937
SUB28_MOD_PM29_L2	159937
Sub Total for MODIS PM	6493544

706-CD-003-001 Day 4 Book A

JD2-31

Benchmark Metadata Composition



SAGE Products

PRODUCT Name	NO. of DB Rows
sag3_m2_gei	29631
sag3_m2_nd	29631
sag3_m2_nt	29631
sag3_m2_o3	29631
sag3_m2_oclo	29631
sag3_m2_tp	29631
sag3_m2_wv	29631
sag3_s1_t1020	30128
sag3_s1_t1550	30128
sag3_s1_t290	30128
sag3_s1_t385	30128
sag3_s1_t440	30128
sag3_s1_t525	30128
sag3_s1_t600	30128

PRODUCT Name	NO. of DB Rows
sag3_s1_t760	30128
sag3_s1_t940	30128
sag3_s2_a1020	30128
sag3_s2_a1550	30128
sag3_s2_a290	30128
sag3_s2_a385	30128
sag3_s2_a440	30128
sag3_s2_a525	30128
sag3_s2_a760	30128
sag3_s2_a940	30128
sag3_s2_gei	30128
sag3_s2_nd	30128
sag3_s2_o3	30128
sag3_s2_tp	30128
sag3_s2_wv	30128
Sub Total for SAGE	870233

706-CD-003-001 Day 4 Book A

JD2-32

Benchmark Metadata Composition



Radar Products

PRODUCT Name	NO. of DB Rows
RADAR1	9580031
RADAR2	9580368
Sub Total for RADAR	19160399

Benchmark Findings



Confirmed:

- Robustness of load procedures for both products
- Ability to load large volumes of data
- Stability of LIBMI and SQS APIs

Will brief at CDR:

- Load times for both products
- Query times for Illustra

Will brief at 3-DR:

- Query times for Sybase

Benchmark Continuation



Timeframe: September through October 1996

Purpose: Benchmark “final” versions of Sybase SQS and Illustra

- Sybase SQS 2.2
- Illustra 3.3
- Beta UDS

Hardware Configuration: Identical to March 1996 benchmark

Data: Current plus 3-D data

Transactions: Same plus 3-D spatial queries

Schema: Release B version modified to handle 3-D spatial queries

Agenda



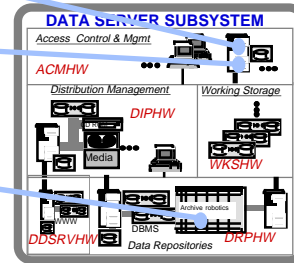
- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- ESDT Design
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics
- Next Steps
- Design Topics

High Level HW/SW Physical Design



Three Processes

- **ScienceDataServer**
 - Process supporting the access to the advertised data type services as well as the implementations of those services.
- **SDSRVAdmin**
 - Supports the SDSRV Administrators and Operators. Primarily access to configuration of the SDSRV and monitor and control of SDSRV processing.
- **COTS Database Engine (Illustra)**
 - Commercial off-the-shelf ORDBMS used for storing and searching metadata.



One Library

- **SDSRVClient Library**
 - Software toolkit library for use by SDSRV Client applications for connecting to and interacting with the SDSRV.

706-CD-003-001 Day 4 Book A

JD2-37

Agenda

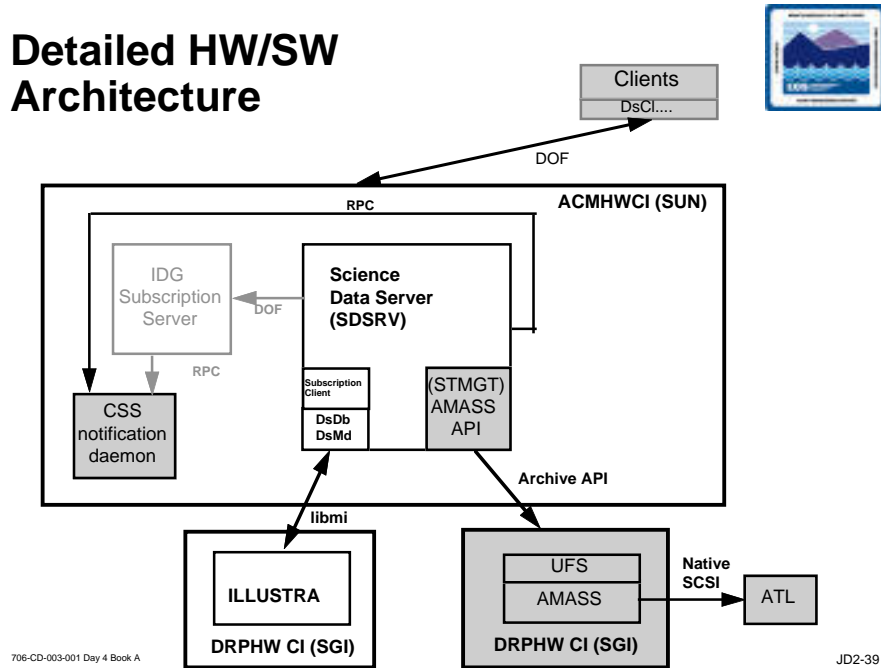


- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- ESDT Design
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics
- Next Steps
- Design Topics

706-CD-003-001 Day 4 Book A

JD2-38

Detailed HW/SW Architecture



Agenda

- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- ESDT Design
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics
- Next Steps
- Design Topics

Transition Assumptions



There is no requirement to run Release A and B SDSRV DBMS Servers in parallel.

Based on Release A DAAC Specific Design Documents, each DAAC will have less than 1GB of metadata to be loaded from Sybase into Illustra.

Migration of Release B V0 data into Illustra during the Release B timeframe is considered a separate transition.

There will be few data type changes or field length changes between the Sybase and Illustra schemas with little “data cleanup” required as in migrations from legacy systems to new systems.

Transition Approach



The transition from Sybase to Illustra will be handled as an “initial database load,” where at the start of Release B, the Sybase DBMS is turned off and the Illustra DBMS is turned on.

Stage 1:

- Construct flat files from Sybase that contain all of the data as of a certain timestamp
 - map Sybase and Illustra schemas
 - develop software to read Sybase and generate flat files that correlate to Illustra tables
- Load data into Illustra using “direct access method”
 - load independent tables in parallel
 - minimize “recovery” time if Illustra goes down during load

Stage 2:

- Construct flat files from Sybase that cover the delta between the initial capture and the cut-off point for the Release B switch
- Load data into Illustra using approach from Stage 1
- Validate that all data from Sybase has been loaded into Illustra

Agenda



- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- ESDT Design
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics

• Next Steps

- Design Topics

706-CD-003-001 Day 4 Book A

JD2-43

Next Steps



Continue Resolution of Design Issues and Topics

- Data Server Detailed Design Meetings (DDM)

Complete Detailed Design

- Release A Retrofit
- Incorporate Updates to Object Model
- Complete Sybase/Illustra Benchmark Prototype
- Complete Schema Design and Update 311-CD-002-005
- Conduct Design Inspections
- Republish 305-CD-024-002

Conduct Delta Detailed Design Review (3DR)

706-CD-003-001 Day 4 Book A

JD2-44

Agenda



- Release B Capabilities
- Software CSCs
- Software CSC Status
- Progress Since IDR-B
- ESDT Design
- Design Trades
- Sybase/Illustra Benchmark Status
- High Level HW/SW Physical Design
- Detailed HW/SW Architecture
- Transition Topics
- Next Steps
- Design Topics

IDR-B Issues



	Issue, Comment, Question, AI, or RID	Code	Requester	Recorder	Comment
2	Want to be able to define data quality level in a subscription.	I	Lyn Oleson	Craig Schillhahn	3DR
6	What is going to be in the metadata? (what is expected to be in the metadata by ECS.	BC	Lyn Oleson	Craig Schillhahn	Completed In DID 311
6	Provide URL for ESDT definitions.	AI	Lyn Oleson	Mike Burnett	Complete
33	How do we group or specify which ESDTs are needed for a specific data product?	I	Lyn Oleson	Craig Schillhahn	The concept of related granule information being packaged with an ESDT granule has been addressed by the granule package. Granule packages are being defined for each ESDT.
34	Where does the doc/presentation discuss the population of Lim/DIM with "granual packages". How can the pieces of a granual package be tracked?	I	Lyn Oleson	Craig Schillhahn	Submitted RID #11. Awaiting Approval
51	Emphasize Update loading in Illustra prototyping. Should interact with DAAC experience over performance expectations. (Lyn offered data.)	I	Lyn Oleson	Steve Marley	Present At CDR
53	Dataserver at the OPS workshop needs to discuss the list of "scripted" subscription actions in response to events that ECS provides upon delivery, and what DAACs are expected to write.	I	Steve Marley	Craig Schillhahn	Conducted Ops Workshop

IDR-B Promises



Number	SDSRV IDR Promises
7	Stability of the Initial ESDTs
7	Operational and Developer Document on ESDTs
8	Establish Application Domain on Session Management
9	Validation of Data Type Service Matrix
10	Bundling of Notifications of Subscriptions
14	Finalize ASF Requirements and Refine the Design
14	Analyze Planning Implications - Nested Production Dependencies
14	Analyze Impact of Run Time Parameters
15	Analyze Feedback from Release A Implementation

RID #5 Meeting Actions



Action Item	Title	Description	Assignment	Contact	Due Date	Comments
4	CDR Functionality Demonstrations	ECS Should include a scenario showing how PDPS, DSS and Client work together. ECS should also present detailed scenarios of virtual product concept using ASTER as an example.	ECS	SDSRV	6-Jun	
5	DAR ID from Japan	Is the DAR ID actually in the metadata?	ECS	SDSRV	Update 4/8	Our design for ASTER DAR ID is to have a product specific metadata attribute for ASTER products. Data Modeling is responsible for identifying product specific metadata for all products.
6	Detailed Subscription Scenario	Provide detailed scenario for how the subscriptions work, e.g., a) how to subscribe to metadata update service, b) belated discovery of corrupted data, and c) manual process of looking up subscribers.	ECS	SDSRV	6-Jun	
7	Notification of Users When Products Change	How to look up all subscriptions going to folks who are PGE-related and all folks who are just science users.	ECS	SDSRV	6-Jun	
8	Different Classes for Data Access	Lyn Oleson suggested DSS might want to consider different classes of users. Need for a flag for data validation and verification prior to access by users?	ECS	SDSRV	6-Jun	
10	Coordination for Data and Browse Ingest	Please provide scenarios that show realistically how metadata and data arrive over time.	ECS	SDSRV	6-Jun	
11	Feedback from Community About "Browse" Definition	Science community has been unable to completely define what and if a browse product is for each of the data products.	ESDIS	Ben Kobler	5-Apr	Ted Meyer is to prepare a paper identifying browse data. ECS is proceeding with our assumption on browse products with 1Mbyte size. DSS doesn't deal with content of package only sends and receives the package. Continues to be open for ESDIS.
14	Temporary URIs	A scenario should be provided describing the steps in a subset request showing the transformation state of the granule.	ECS/CDR	SDSRV	6-Jun	
17	Coordination for Metadata Problem Reporting	We need to be sure that CIDM knows how to handle the DSS metadata problem report and that there is no confusion with Trouble Tickets. Need to assess the role of metadata problem report in the system context.	ECS	SDSRV	6-Jun	

RID #5 Meeting Actions (cont.)



Action Item	Title	Description	Assignment	Contact	Due Date	Comments
18	Description of DAR Requirements	What subsystem at CDR will describe the DAR scenario? Associated with the DAR, what "up front" capabilities will be available in the client GUI tool?	CDR	SDSRV	6-Jun	
19	How Algorithms are Made Available	A scenario for acquiring the Science SW Algorithm Programs from the archive should be included in the documentation. Includes several issues.	ECS	SDSRV	6-Jun	
20	Description of How New Product Versions are Handled	State explicitly what versioning is and what the requirements are for it. Show how the system is designed to handle it.	CDR	SDSRV	6-Jun	
23	How Many Data Granules Represented in DB Benchmark?	How is the Benchmark going to simulate the varying data types representative of Releases A and B? Please include the number of each different type in the study (how many rows per group?)	CDR	SDSRV	5-Apr	Approximately 46 Million
24	How Does Subscription Processing Impact Database Performance?	There was a question about the distribution of the 100 users being simulated. Several suggestions were made.	ECS	SDSRV	5-Apr	There is not a subscription simulator but the subscription load could be considered part of the 100 concurrent users. SDSRV doesn't distinguish between actions being fulfilled as actual requests vs. on behalf of clients from the subscription server.
26	How do Priorities in Data server Allow Management of Limited Resources?	Chris Harris requested that we determine what the dynamics of changing priorities are and how many different types of classes can you have at a given time? How many flavors of users are there?	ECS	SDSRV	6-Jun	
27	How is Priority Information in DARE and in On-Demand Processing Passed to Processing?	DAR processing scenarios were discussed; a controversial point concerned how the user profile was passed to the SDSRV. Oleon stressed how different levels of users may submit a DAR.	ECS	SDSRV	6-Jun	
29	QA for On-Demand Processing and Re-Generation of Higher Level Products After Data Restoration	Check on QA plans both in regular processing and in the on-demand processing.	ECS	SDSRV	6-Jun	

706-CD-003-001 Day 4 Book A

JD2-49

Ops Workshop Action Items



SDSRV Action Items Number	Title	Description	Status	Due Date	Document Trace
Action Item #9	Verify Metadata	Need clarification on "verify metadata" step. Are there unnecessary redundancies?	Closed		
Action Item #10	View of Inventory	Desire to control the view of the inventory on a granule/file level. Explore a public vs. private (restricted) advertising of data availability as a means of restricting access to a targeted granule/file.	Open		
Action Item #11	EOS Products	Need to identify file dependencies on EOS products & VO migrated data.	Open		
Action Item #12	Partial Failure	Is requestor notified of a partial failure of a request?	For store and retrieval requests: No, there are no partial successes for store or retrieval operations.		
Action Item #51	Submitted Requests	For submitted requests by User Services on a user's behalf, user services will let the user know how big their request is.	Reject -User Services		
Action Item #54	Sybase to Illustra Switchover	What will happen if during the switchover from Sybase to Illustra, it is discovered at a later date that the Illustra DB is corrupted to the point that reverting back to Sybase without lost data is not possible?	Accept		
Action Item #55	Sybase to Illustra Switchover	Consider the number of granules as well as volume for the Sybase to Illustra and VO transition include in the white paper.	Open		
Action Item #57	Fix Scenario Disconnects	Fix any disconnects between this scenario and the subscription event error/resolution scenarios.	Open		
Action Item #73	Order processing Cancellation	At what point in processing is it too late to cancel an order?	Closed		
Action Item #78	Data Access Restrictions	Are data access restrictions reflected in the metadata model? If not, they should be.	Reject		
Action Item #85	Group Subscriptions	Need a scenario to show how group subscriptions will be handled. (How do you do a batch notify of a problem?)	Rejected - IDG Function		

706-CD-003-001 Day 4 Book A

JD2-50